Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story
Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles
Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor
To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives
Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking
We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

Introducing the Intel® Pentium® 4 Processor

Christie Rice
Processor Marketing Manager
Desktop Platform Group
Intel Corporation

Overview

Bandwidth-hungry Internet and advanced applications are creating opportunities to think in new ways. Intel has risen to the challenge by designing a platform that meets today's high-end performance needs while providing headroom for technologies and usages that are still under development today.

With the Intel® Pentium® 4 processor, Intel is delivering world-class performance across both existing and emerging applications and usage models, delivering performance headroom and scalability for the future. In addition, the Intel Pentium 4 processor is an ideal solution for sophisticated end users working with complex Internet, imaging, video, speech, 3D, and multimedia applications.

The Intel® NetBurst™ microarchitecture is the foundation for the Intel Pentium 4 processor and provides a number of new and improved architectural features that take desktop performance to a new level. These new features include a 400-MHz system bus, hyper-pipelined technology, advanced dynamic execution, rapid execution engine, advanced transfer cache, execution trace cache, and Streaming SIMD (Single Instruction, Multiple Data) Extensions 2 (SSE2).

400-MHz System Bus

The Intel Pentium 4 processor features a 400-MHz system bus (the 100-MHz clock is quad-pumped) that provides 3X bandwidth over the 133-MHz system bus in the Intel® Pentium® III processor (1.06 GB/s). With 3.2 GB/s of system bandwidth, the Intel Pentium 4 processor delivers the highest bandwidth desktop bus currently in the industry.

Hyper-pipelined Technology

With the Intel Pentium 4 processor, Intel has doubled the pipeline depth to 20 stages, enabling a higher clock frequency. The additional pipeline stages establish a new baseline for processor speed, delivering ≥1.40 GHz at launch on our 0.18 micron process. This higher core frequency significantly increases processor performance and frequency capability and provides the scalability needed for future applications.

Advanced Dynamic Execution

The Advanced Dynamic Execution engine is a very deep, out-of-order speculative execution engine that keeps the execution units executing instructions. It does so by providing a very large window of instructions from which the execution units can choose. The large out-of-order instruction window allows the processor to significantly reduce stalls that can occur while instructions are waiting for dependencies to resolve. One of the more common forms of stalls is waiting for data to be loaded from memory on a cache miss. This aspect is very important in high-frequency designs, as the latency to main memory increases relative to the core frequency. The NetBurst microarchitecture can have up to 126 instructions in this window (in flight) vs. the previous P6 microarchitecture's much smaller window of 42 instructions.

The Advanced Dynamic Execution engine also delivers an enhanced branch prediction capability that allows the Pentium 4 processor to be more accurate in predicting program branches. This has the net effect of reducing the number of branch mispredictions by about 33 percent over the P6 microarchitecture's branch prediction capability. It does this by implementing a 4-KB branch target buffer that stores more detail on the history of past branches, as well as by implementing a more advanced branch prediction algorithm. This enhanced branch prediction capability is one of the key design elements that reduce the overall sensitivity of the NetBurst microarchitecture to the branch misprediction penalty.

Rapid Execution Engine

The two Arithmetic Logic Units (ALUs) in the Intel Pentium 4 processor run at twice the core frequency of the processor. This makes it possible to execute basic integer instructions (such as add, subtract, logical AND, and logical OR) in half a clock cycle, with higher execution throughput and reduced latency of execution. With the 1.40-GHz Intel Pentium 4 processor, each of the ALUs is running at 2.80 GHz, increasing performance on integer-based applications.

Revolutionary Cache Subsystem

In order to increase performance and scalability, the Intel Pentium 4 processor features an innovative new cache subsystem designed to optimize data transfer to the core. An execution trace cache stores 12K decoded instructions in the order of program flow instead of predecoded instructions that cannot take code branches into consideration. The execution trace cache removes the decoder from the main instruction loop and results in a higher performance, more efficient level 1 instruction cache.

The Intel Pentium 4 processor also includes a level 2 Advanced Transfer Cache (ATC). While still only 256 KB in size, this ATC improves the data transfer rate between the on-die level 2 cache and the processor core to 44.8 GB/s at 1.40 GHz, compared with 16 GB/s on a 1-GHz Intel Pentium III processor. The level 2 Advanced Transfer Cache is able to clock 256 bits (32 bytes) of data into and out of the cache on every clock cycle, unlike previous microarchitectures. The overall gain with the new cache subsystem is that the transfer rate between the cache subsystem and the processor core is optimized over previous subsystems in both bandwidth and latency. The enhanced cache subsystem delivers increased performance and response on a wide variety of applications.

Streaming SIMD Extensions 2 (SSE2)

To make the SIMD instruction set even more powerful, the Intel Pentium 4 processor provides 144 new performance improving instructions, including 128 bit SIMD double precision floating point, 128-bit SIMD integer, and improved cache and memory management instructions.

The SSE2 extends Intel's MMX™ technology to 128 bits and supports packed integer operations. While the extended width of the operation used to be 64 bits, these new instructions double the SIMD integer bandwidth over SSE/MMX technology. This accelerates a broad range of applications, including video, speech, and image and photo processing.

The new 64-bit adds/subtracts and 32x32 unsigned multiply provide significant enhancements to encryption operations as well. As we move into the future, encryption/decryption capabilities will be more important in driving a secure e-Business infrastructure for the connected world.

The 128-bit SIMD double precision floating-point delivers the capability to execute two 64-bit double precision floating point instructions at once, doubling the performance capability. In addition, it offers a full set of SIMD double precision floating-point operations, and additional operations that convert between double and single precision. This precision floating point results in the acceleration of content creation, financial, engineering, and scientific applications.

Balanced Platform Solution

As part of a complete platform solution, the Intel Pentium 4 processor was designed in tandem with the Intel® 850 chipset to create a powerful new platform for high-performance users. The 400-MHz system bus in the processor is balanced by dual RDRAM* memory channels in the 850 chipset that operate in lock-step to deliver 3.2 GB/s of memory bandwidth. Coupled with more efficient protocols and the 400-MHz system bus, the Intel Pentium 4 processor and Intel 850 chipset deliver three times the bandwidth of platforms based on high-performance Intel Pentium III processors. The increased bandwidth enables faster memory acquisitions, which increase performance on any application requiring intensive memory accesses such as many 3D and video applications.

Summary

The Intel Pentium 4 processor was developed for today's high-end applications and fast-evolving Internet technologies. With today's increased power and performance demands, the Intel Pentium 4 processor offers an ideal platform solution for 3D streaming content on the Web or interactive gaming.

As the foundation for the next generation of Intel® IA-32 processors, the Intel NetBurst microarchitecture offers exceptional performance advantages and architectural innovation. Designed and optimized to provide incredible performance over a range of applications and uses, the Pentium 4 processor delivers breakthrough performance across all applications where end users will truly be able to experience and appreciate the performance.

More Info

For more information on the Intel® Pentium® 4 processor, visit the Intel Pentium 4 processor section of the Intel Web site.

The Intel Pentium 4 processor presentation (PDF, 537K) delivered at the Intel Developer Forum Conference, Spring 2000, is also available from the Intel Pentium 4 processor section of the Intel Web site as is the presentation from the Fall 2000 Conference (PDF, 811K).

Author Bio

Christie Rice is the processor marketing manager for Intel's Desktop Platform Group. During her four years with Intel, Christie has been part of the marketing team for both the Intel Pentium III processor and the new Intel Pentium 4 processor. -- Christie began her Intel career with the Flash memory group, where she earned industry recognition for enabling the industry to store code and data in Flash components. Prior to joining Intel, she worked for 14 years in circuit board and system design and marketing. -- A graduate of New Mexico State University, Christie has a B.S. in electrical engineering. She also has an M.S. in electrical engineering from Southern Methodist University and an M.B.A. from the University of Dallas.

Columns

From the Editor
Donna Loveland
Managing Editor
Intel Developer Update Magazine
Intel Corporation

Column

Developers and consumers alike have been asking what Intel would do to follow up on the Pentium® III processor. Our cover story gives you a first look at the answer with a tour of the microarchitecture of the Intel® Pentium® 4 processor. Exciting as this new processor development is, Intel also has plenty going on in the areas of desktop, software, networking and communications, and initiatives and technologies. Here's the latest:

Intel® Pentium® 4 Processor—An Overview—cover story—As the foundation for the next generation of Intel® IA-32 processors, the Intel Pentium 4 processor supports high-end applications and fast-evolving Internet technologies. The Intel® NetBurst™ microarchitecture provides a number of new and improved architectural features that take desktop performance to a new level.

DSL on CNR for the Home Server PC—The "home server" PC provides a central location through which all the networked PCs within a home can communicate with the outside world. The CNR (Communication and Networking Riser) specification provides a flexible and cost-effective way to implement DSL, V.90 modem, and home phoneline networking in a home server PC.

Building Trust and Privacy into Open PC Systems—The PC platform and Internet infrastructure enable e-Business -- and provide opportunities for hackers and malicious individuals. With Version 0.9 of its security specification, the Trusted Computing Platform Alliance (TCPA) has achieved solid results in its work to build a foundation for improved trust.

Security and Confidentiality for Intel® Early Access Services—Intel offers software developers free access to new technology through the Intel Early Access Services Web site. Recent advances in security technology combat unauthorized access and activity, making it possible for Intel and other companies to offer shared access to remote systems.

Reference Designs for Network Attached Storage—The expansion of the Internet, pervasive networking, and the emergence of richer file types are all driving demand for additional network data storage capacity. Dedicated Network Attached Storage (NAS) appliances provide a simple, affordable storage solution for high-availability networks; Intel offers two free reference designs.

CNR-Enabled Wireless—Bluetooth and 802.11b are potential implementations for desktop personal computers, and the Communications and Networking Riser (CNR) specification provides a flexible and cost-effective way to implement wireless, DSL, V.90 modem, and home phoneline networking in desktop PCs.

Be sure to use the links in the "More Info" section of each article to get details on these significant topics.

Enjoy.

Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Platform Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held technical and marketing positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in Expository Writing from the University of Iowa.

Departments

Desktop

CNR-Enabled Wireless
Deepti Gupta
Product Marketing Engineer
DPG Boards & System Marketing
Intel Corporation

Overview

In the wireless communications industry Bluetooth* and 802.11b are well-known technologies for devices such as cellular and mobile. These technologies are not as well known as potential implementations for desktop personal computers (PCs). However, the growth of Bluetooth-enabled desktop PCs is expected to be explosive, from 2 million in 2001 to 81 million by the year 2005 (Cahners In-stat Group, July 2000, see Chart1).

Projected Bluetooth-Enabled Devices (Numbers in Millions)
Cahners In-Stat Group
July 2000

BT-Enabled Device	1999	2000	2001	2002	2003	2004	2005
Desktop PCs			1.954	8.978	24.389	51.593	81.078
Notebooks		0.112	1.028	6.278	17.544	32.386	52.240
Inkjet Printers			2.040	6.320	13.650	27.810	53.775
B&W Laser Printers			0.053	0.158	0.362	0.642	1.150
Cordless Phones			0.806	2.528	5.528	12.795	22.563
PBX Systems			0.009	0.032	0.100	0.180	0.350
Headsets		0.002	0.154	1.493	9.892	25.078	58.978
Handheld PC & PDA			0.047	0.306	1.255	3.424	6.570
Digital Cameras			0.075	0.285	0.825	1.722	3.105

Communications and Networking Riser, or CNR, is an open industry specification for desktop PCs. It enables technologies such as Bluetooth, audio, modem, LANs (local area networks), HPNA (Home Phoneline Networking Alliance), 802.11b (wireless) technologies; and broadband solutions, such as digital subscriber line (DSL) and cable modem. CNR is one of the most advantageous Bluetooth and 802.11b implementations for cost, flexibility, homologation, and real-estate efficiency in desktop PCs.

CNR supports a variety of interfaces (see Figure 1) for these technologies. These interfaces include the Audio Codec ’97, system management bus (SMBus), universal serial bus (USB), and either Intel’s LAN-connect interface (LCI) or the media-independent interface (MII). Note that USB signals must be implemented to enable 802.11b and Bluetooth on a CNR.

The Communications and Networking Riser Interface

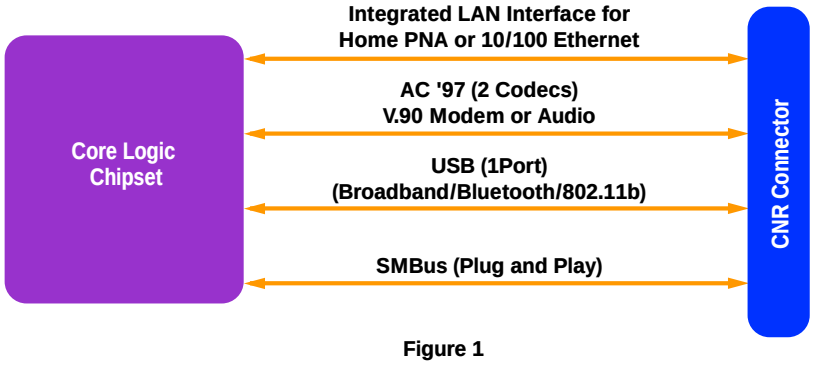

Figure 1

Internal Implementation

CNR is an inside-the-box implementation. The Bluetooth CNR architecture (see Figure 2) consists of three major parts: a baseband, a link controller, a radio, and an antenna. The Bluetooth Module communicates with the system's host controller using a high-speed USB interface. The module appears as a USB slave device and thus requires no PCI (peripheral component interconnect) resources.

The 802.11b CNR architecture (see Figure 3) similarly includes a baseband, media-access controller (MAC), and a radio.

Figure 2. Bluetooth on CNR architecture for a desktop PC

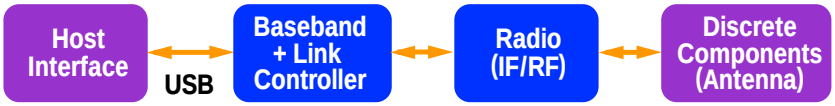

Figure 3. 802.11b on CNR architecture for a desktop PC

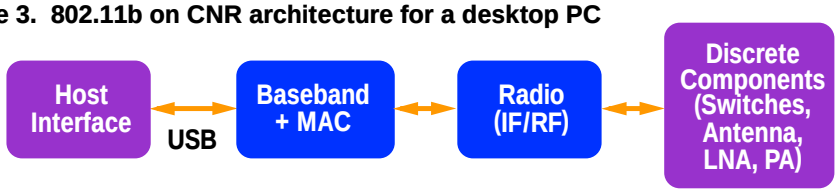

Advantages of CNR

Implementing Bluetooth and 802.11b on CNR offers a number of advantages:

- The multiple interfaces supported by CNR allow flexible combinations of technologies to coexist on a single internal CNR card, based on the requirements of the specific platform.
- System integrators can use same PC platform in different market segments (such as consumer and business segments) by offering different CNR solutions. This reduces board SKU (stockkeeping unit) count, management, and inventory control, as well as customer-support issues, all of which represent cost savings.
- The SMBus interface provides CNR with plug-and-play functionality. This means that CNR helps developers get Windows* Hardware Quality Lab (WHQL) certification for their desktop designs.
- The mechanical design of the CNR card and the location of the connector allow CNR to share the slot location for a PCI-IO (input/output). In other words, CNR does not consume a PCI slot when the CNR slot on the mainboard or motherboard is not being used.
- Issues of noise, electrical interference, and real-estate requirements on the motherboard are reduced.
- Homologation issues are resolved, and FCC certification carries across multiple platforms.

Summary

The CNR specification provides a flexible and cost-effective way to implement wireless, DSL, V.90 modem, and home phone line networking in desktop PCs. CNR incorporates multiple interfaces on a single interface connector. As a result, several communication and networking building blocks can be easily incorporated onto a single card. This in turn provides OEMs, manufacturers, and system integrators with the flexibility to integrate PC functionality as a value-added feature and at a price point that is potentially much lower than what can be achieved using traditional PCI expansion cards.

With CNR, users don't have to look for external Bluetooth or 802.11b solutions for their desktop PCs. CNR enables these two new and exciting technologies as internal PC implementations.

Projections show that by 2005 there will be over 81 million BT-enabled desktop PCs. Developers should begin implementing USB signals on the CNR now, in order to support both current and future BT, 802.11b, and other broadband implementations.

More Info

Some aspects of CNR (including supported interfaces) are discussed in the IDU article titled "CNR Card Offers Motherboard Expansion." This article, which appeared in the May 2000 issue, describes CNR flexibility and how CNR can be used to implement internal modems, wired and wireless LAN, audio, etc. The article also discusses putting the USB signals on the CNR card (which CNR supports) to enable new technologies such as Bluetooth and 802.11b.

Bluetooth technologies are described in detail at the Bluetooth Web site. The 802.11b specification is also available online.

Author Bio

Deepti Gupta, a product marketing engineer for wireless technologies, has been with Intel for six years. She has worked as a technical marketing engineer for motherboards and systems, and for concept platforms. Deepti has received several Intel awards including recognition for her work with WHQL program activities. Deepti received her B.S.E.E. from the University of Washington.

DSL on CNR for the Home Server PC

Brad A. Barmore
PC Audio & Communications Architect
OEM Platform Solutions Division
Desktop Platform Group

Overview

The "home server" PC provides a central location through which all of the networked PCs within the home can communicate to the outside world. To serve as a server, this PC must support Wide Area Network (WAN) technologies such as DSL (Digital Subscriber Line), cable modem, or V.90, in addition to a Local Area Network (LAN) technology, such as home phoneline or wireless networking.

As the demand for DSL continues to grow, one possible home server configuration supports the combination of DSL and V.90 technologies with home phoneline networking. Each of these technologies uses the telephone line RJ-11 jack to communicate. When all three of these technologies are integrated onto a single card, the OEM or system integrator can ship a single product that supports a DSL or V.90 WAN, plus a home phoneline LAN, without the need to provide additional telephone cables, RJ-11 "Y" connectors, and filters. The integration of these technologies on a single card reduces the potential support costs associated with implementing the same technologies on individual cards.

The CNR (Communication and Networking Riser) specification provides a flexible and cost- effective way to implement DSL, V.90 modem, and home phoneline networking in a home server PC. CNR incorporates multiple interfaces on a single connector, and as a result, several communication and networking building blocks can be easily incorporated onto a single card. This in turn provides OEMs, manufacturers, and system integrators with the flexibility to integrate home server PC functionality as a value-add, at a price point that is potentially much lower than what can be achieved using traditional PCI expansion cards.

Multiple Interfaces

As Figure 1 shows, CNR technology flexibly supports a variety of WAN and LAN technologies through its various interfaces. The interfaces included on CNR include the Audio Codec '97 (AC '97), System Management Bus (SMBus), Universal Serial Bus (USB), and either Intel's LAN Connect Interface (LCI) or the Media Independent Interface (MII).

These interfaces allow flexible combinations of audio, DSL modems, cable modems, V.90 modems, 10/100 Ethernet LAN, home phoneline networking, or wireless networking all to coexist on a single card, based on the requirements of specific platforms. In today's PCs the controller side of each of these interfaces is typically integrated into the core logic chipset.

As shown in Figure 1, the V.90 analog modem is supported through the AC '97 Interface. Home phoneline networking is supported through the LCI Interface on Type A CNR connectors or the MII Interface on Type B connectors. Though DSL does not have a dedicated interface, it is easily supported through the USB Interface.

The Communications and Networking Riser Interface

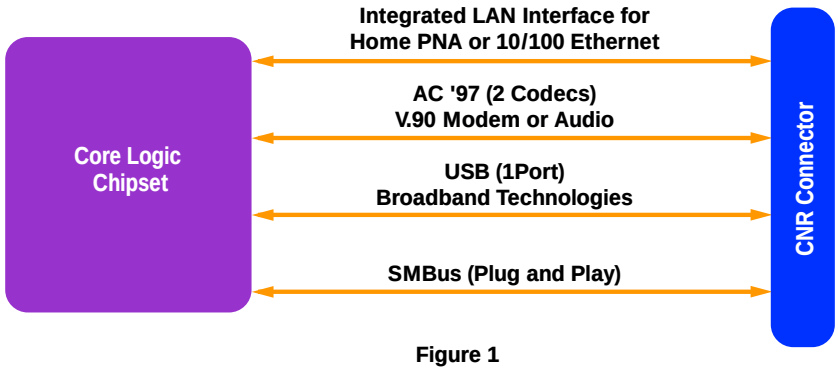

Figure 1

DSL Modem on CNR

CNR can easily support a DSL modem using the existing USB interface and USB-based DSL silicon. OEMs, system integrators, and CNR manufacturers do not need to create a new interface or wait for the silicon development on a new interface.

As shown in Figure 2, DSL modems typically use an architecture that includes a bridge device between a common bus (e.g., PCI, MII, or USB) and a Digital Signal Processor (DSP). The DSP then communicates through a separate interface to an Analog Front End (AFE). The final blocks in this architecture include the Line Driver and Receive Filters.

DSL on CNR Architecture for a Home Server PC

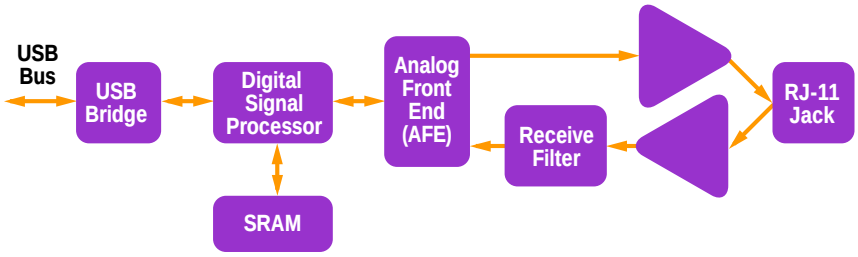

Figure 2

In the DSL on CNR architecture, the data to be transmitted over the DSL modem is transferred to the PC's USB controller. It is then formatted and sent to the USB bridge device on the CNR, where the data is extracted and sent to the DSP. The DSP then takes this data and applies the appropriate encoding and modulation techniques and transmits it to the AFE. The AFE device essentially consists of a Digital to Analog Converter (DAC). This DAC then converts the encoded and modulated data from the DSP and turns it into an analog signal which is finally sent to the Line Driver for amplification and out to the telephone line and to the telephone company Central Office.

Data that is received via the telephone line is passed through the receive amplifier and receive filter to the AFE. The AFE then digitizes the encoded and modulated data using an Analog to Digital Converter (ADC). The digitized data is then passed into the DSP where it is demodulated and decoded and passed to the USB bridge. Finally, the USB bridge device formats the data for transmission over the USB Interface and to the PC's USB controller where it is received and sent back to the system memory for use by the PC.

Summary

CNR provides OEMs and system integrators with a flexible and cost-effective solution for implementing new communications and networking technologies. As DSL modems and home networking become more popular in the consumer market segment, the “home server” PC, integrating DSL, V.90, and home phoneline technologies, provide a way for networked home PCs to share a DSL connection, all through the use of a single RJ-11 connector.

CNR technology provides a more cost-effective integrated solution by providing an inexpensive method for implementing DSL, V.90, and home phoneline networking technologies. Implementing server functionality is simply a matter of integrating a single riser card that provides a single RJ-11 jack, eliminating the requirement for additional phone cable, Y connectors, and specialized filters. In addition, the solution uses industry standard silicon building blocks from a variety of manufacturers.

While CNR technology provides a lower cost, time-to-market solution for OEMs, CNR manufacturers and system integrators, it can also help simplify the in-home deployment of shared broadband connections, providing another great way to enhance the PC user experience.

More Info

The Intel® CNR Web site includes press releases, industry links, and downloadable information for OEMs, system integrators, and CNR manufacturers. Visit the Web site to download version 1.1 of the CNR Specification and version 1.0 of the CNR System Design Guide in Adobe Acrobat* format.

For an overview of the benefits of CNR for developers, see the article CNR Card Offers Motherboard Expansion by K.L. Yeung in the May 2000 *Intel Developer Update*.

For information on home phoneline networking, visit the Home Phoneline Alliance Web site.

Author Bio

Brad Barmore is a PC audio and communications architect. His industry contributions include authorship of the Communications and Networking Riser Specification and co-authorship of the Audio/Modem Riser Specifications. In addition, Brad continues to be the primary architect of the audio and modem subsystems for OPSD's desktop PC motherboards. He holds patents in audio circuitry and has multiple patents pending in the area of PC riser technologies. He graduated from Washington State University in 1986 with a B.S. in electrical engineering.

Initiatives and Technologies

Building Trust and Privacy into Open PC Systems

David Grawrock
Security Architect
TCPA Technical Committee Member
Desktop Architecture Lab

Overview

All business and commerce depend on trust, and this is especially true of e-Business. The openness of the PC platform and the open infrastructure of the Internet provide the essential flexibility to enable widespread adoption of a variety of innovative e-Business applications, including innovative business-to-consumer (B2C) and business-to-business (B2B) systems. While these systems have the power to generate billions in revenue, it is also true that "threat follows value." The same quality of openness that enables e-Business also provides opportunities for hackers and malicious individuals. Because e-Business transactions occur on the PC, enhancing trust in the computing platform itself has become an issue of real importance for OEMs, hardware vendors, and software developers.

In the spring of 1999, the Trusted Computing Platform Alliance (TCPA) was chartered to encourage widespread industry participation in the development and adoption of a new open platform-based security specification. The goal of this ongoing effort is to build a solid foundation for improved trust in computing platforms over time, with the initial version of the specification focusing on the PC. TCPA participants have further agreed that the specification for the trusted computing PC platform should focus on two areas, enhancing security and protecting privacy.

This effort has achieved solid results. Version 0.9 of the TCPA specification was publicly released at the Intel Developer Forum (IDF) Fall 2000 in August. The well-attended TCPA track at IDF featured an introduction to the Trusted Client concept, followed by detailed technical sessions on protected storage, ownership of the trusted client module (TPM), authenticated anonymity, integrity metrics/authenticated boot, run-time use of TCPA features, and TCPA compliance issues. A Trusted Clients Panel followed the technical presentations, and IDF featured demos of TPM solutions presented by TCPA member companies.

The TCPA has grown to more than 140 member companies and it is working to complete version 1.0 of the specification before the end of 2000. Joining the TCPA and understanding the specification is advisable for anyone in the industry who has an interest in building trust into computing platforms while protecting privacy for users.

Why We Need TCPA

Today's PCs offer some level of trust to users, and e-Business as it exists today depends on the trust available from modern PC systems. However, as PCs evolve and as e-Business becomes more sophisticated, we are also seeing the evolution of requirements for additional security. With the growing demand for trusted computing, the ability to protect the integrity of a PC through software alone is reaching its limitations. In order to ensure the trustworthiness of an OS in an open system, one must start from a "root of trust" that begins from the earliest point of booting the system. Future operating systems should be able to provide enhanced trust, and their use of the TCPA can become the basis for such improvements.

Platform-based Trust

To overcome this limitation, the TCPA specification defines a foundation of platform-based trust that is extendable across systems and networks for multiple users. The goal of the TCPA, through the TCPA specification, is to make trust just as much a part of computing platforms as memory and graphics. Every TCPA-compliant platform will integrate this trusted functionality, and every piece of software on the system will be able to use it.

The TCPA specification is designed to enhance the deployment, use, and manageability of platform security elements. In addition, it provides the industry with pervasive standards that are designed to be both cost-effective and exportable. The specification provides the foundation for a baseline security standard that can then be extended by individual vendors. The TCPA specification is also designed to complement existing capabilities, including the X.509 standard for digital certificates, IPSEC (Internet Protocol Security Protocol), IKE (Internet Key Exchange), VPN (Virtual Private Network), PKI (Public Key Infrastructure), PC/SC Specification for smart cards, biometrics, S/MIME (Secure Multipurpose Internet Mail Extensions), SSL (Secure Sockets Layer), and SET (Secure Electronic Transaction).

Definitions

A trusted platform is defined by TCPA as "a platform that behaves the way it is expected to behave for an intended purpose." The basis for this trust is a declaration by a known authority that the platform can be trusted for an intended purpose. In other words, someone whom the owner trusts says with certainty that a platform is exactly what it says it is, and that the owner can trust the platform.

The TCPA specification defines a separate "module," called the Trusted Computing Subsystem (TPS), that can be trusted to the same degree by as many entities as possible. It consists of two building blocks:

- *The Trusted Platform Module (TPM)*, defined as a supplier of protected functionality and shielded locations. Protected functionality resides within the TPM to guarantee authorization and correct use. Shielded locations contain information that can be accessed only by protected functions.
- *BIOS software* to perform "integrity metrics." In conjunction with the TPM, this provides the measurement root of trust.

Through the TPS, the TCPA specification creates a "chain of trust" which is built from the system's first operation through the OS.

Authenticated Boot

The TCPA-authenticated boot process depicted in Figure 1 provides a good example of transitive trust. During authenticated boot the TCPA-compliant platform measures its boot process and stores the results of that measurement in the TPM. Integrity metrics can be derived from the BIOS, boot-loader, OS loader, and the OS security policy, using cryptographic hashing to extend transitive trust from the BIOS to other areas of the platform. There is no attempt to stop the boot process nor to change the process flow. The stored measurements can be used by challengers to determine what boot process occurred. The measurement can be reported multiple times. The benefit is that challengers of the platform can trust that the measurement, storage, and reporting of integrity metrics is done in a secure manner.

The measurement process takes a cryptographic hash of the entity to measure. For authenticated boot, this starts with the hashing of the BIOS. The resulting hash is stored in the TPM. When challenged, the TPM digitally signs the stored integrity metric and returns it to the challenger. The challenger bases its trust in the remote platform on the credential associated with the key used to sign the integrity metrics.

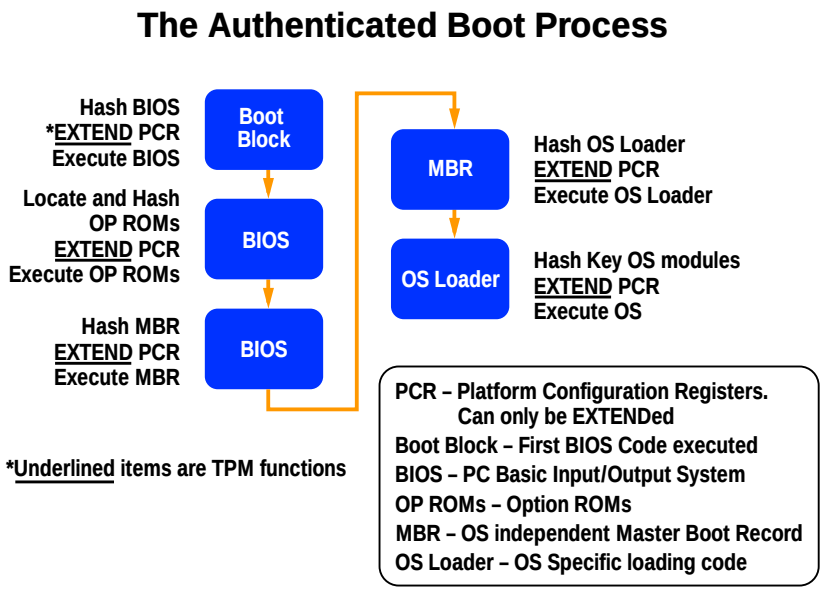

Figure 1

The trust that a challenger places on a remote platform is based on transitive trust. The transitive process starts from the root of measurement trust, which measures the next application, stores the measurement, and then passes control on to the next application (e.g., from the BIOS to the option ROMs to the operating system loader, and so on). If the measurement is done correctly, the challenger knows what the chain of applications is. If a "bad" application exists, the bad application cannot perform the right measurement. In addition, the previous application would have measured the bad application and reported it. The bad application can still run; it just cannot hide its presence. This preserves the openness of the system while enabling a chain of trust to be built from the trust root. The challenger can make a decision regarding the trustworthiness of the remote platform for its intended purpose.

Retaining Privacy

Individuals and businesses have confidential data to protect. This makes privacy an important element of any trusted system. In this context, "privacy" can be defined as a way to prevent others from gaining access to protected information without the informed consent of its owners.

The TCPA specification includes specific steps to preserve privacy, while enhancing trust. The TCPA specification defines mechanisms, not policies. The objective is to enable the reporting of integrity metrics by the platform without restricting the choices available to the system owner. For example, the owner maintains ultimate control over all private information and must proactively "opt-in" to use the TCPA subsystem.

To further enhance privacy, the TCPA specification lets the system owner create multiple anonymous identities to enhance personal security and remove avenues for identity cross-correlation. The TCPA specification also provides an identity protocol to preserve anonymity. These privacy features help to preserve the quality of openness in the PC platform and can foster a higher level of user confidence in TCPA-enabled platforms.

Demos at IDF

The TCPA track at IDF Fall 2000 included a variety of software and hardware demos from TCPA members including Compaq Corporation; Fujitsu Siemens Computers GmbH; Hewlett-Packard; Infineer, Inc.; Infineon Technologies AG; Intel Corporation; Microsoft Corporation; National Semiconductor Corporation; Phoenix Technologies, Ltd.; ST Microelectronics; Schlumberger, Ltd.; and Wave Systems.

The demos included a Network Gatekeeper proof-of-concept demonstration that featured the basic functionality of a TPM with integrity metrics, protected storage, and the challenge-and-response mechanism. The demonstration featured a server page that allowed access to a protected resource only when the client provides proof of the client's status through the authenticated boot process.

As an active member of the TCPA, Intel encourages the wide industry support and adoption of TCPA-based solutions

and products. When adopted, TCPA solutions will enable users to enhance security while creating opportunities for the emergence of new e-Business models and new ways of protecting systems through the use of TCPA capabilities.

Summary

The Trusted Computing Platform Alliance (TCPA) has developed and released a specification that is designed to simplify and accelerate the deployment, use, and manageability of security capabilities on computers. A TCPA-enabled system provides a cost-effective and standardized way to embed security functionality in a platform. This enables improved levels of security to become ubiquitous, while encouraging the development of applications that use security. A TCPA subsystem also improves the control of the way data is accessed.

Traditionally, access has depended upon authorization or authentication. Under TCPA, access can be linked to the state of the software in the platform. If rogue software, such as a virus, is introduced into a platform, access can be denied, based on the change in the software state of the platform. Ubiquitous platform-based security encourages the development and use of security services. PKI-related security processes, such as digital signature and key exchange, are protected through the secure TCPA subsystem.

The release of the TCPA specification provides an excellent opportunity for software developers and independent hardware vendors to join and participate in the TCPA. Hardware vendors will need to understand the specification in order to enable TCPA security in their products. Software developers should optimize their applications to benefit from TCPA services.

More Info

The TCPA Web site contains detailed information about the Trusted Computing Platform Alliance, including the IDF Fall 2000 presentations, upcoming events, white papers, press announcements, and details on how new members can participate. The Web site provides TCPA members with access to the recently released v0.9 specification, outlines technical issues, and enables TCPA members to provide important feedback.

Author Bio

Now in his second year with Intel, David Grawrock is a security architect and TCPA program manager in the Intel Desktop Architecture Lab. As a member of the TCPA Technical Committee, he has been directly involved in the development of the TCPA specification. His presentations at IDF Fall 2000 included Trusted Computing Module Ownership and TCPA Compliance. Prior to joining Intel, David was the architect of the *Norton Your Eyes Only* desktop security application, from Symantec Corporation. He holds five patents in the area of cryptography and computer security. He is a member of the International Association of Cryptographic Research (IACR).

Networking and Communications

Reference Designs for Network Attached Storage

David Hillyard
Communications Platform Manager
Embedded Intel Architecture Division
Intel Corporation

Overview

The dramatic expansion of the Internet, pervasive networking, and the emergence of richer file types are all driving demand for additional network data storage capacity. Dedicated Network Attached Storage (NAS) appliances provide a simple and affordable storage solution for high-availability networks.

Network Attached Storage Appliances

The traditional way to add storage capacity to a network is to add disk drives to servers or install a file server on the network. Both of these approaches are complex and expensive, and they can interrupt data availability because the server must be offline during the installation process. Compared to a PC-based file server, NAS appliances cost much less to buy and maintain. NAS appliances have simplified installations with self-configuration, permitting the network to remain up and running. Adding more storage is as simple as plugging in another NAS appliance.

NAS appliances have a number of other advantages. They support cross-platform file sharing in heterogeneous network environments. Their versatile usage model includes file sharing, online storage, and data backup. In addition, the compact form factor of NAS appliances makes them a practical, transportable storage solution.

NAS appliances can also improve network data flow and performance. When used to cache frequently accessed files, NAS servers can remove potential bottlenecks from the network's general-purpose server, especially when the network server must support large numbers of requests from multiple users. NAS appliances used for file caching can also be located in close proximity to the users who need the data, improving throughput time.

Journaling File System

A common method of maintaining performance as well as reliability within NAS appliances is through the use of a Journaling File System (JFS). In a JFS fast file system, recovery and restart is made possible through logging of file activity within a byte-level file system. In the event of a system failure such as a system crash, no file system transactions are left in incomplete states. The result is that the file system can be restored in a matter of seconds, versus minutes, as would be the case with non-journaling file systems.

There are two primary components within a JFS. The log file is responsible for logging operations such as file creation, linking, making directories and nodes, removing files, renaming, setting ACL, and writing files among others. A second component of a JFS is the transaction manager. It is responsible for providing central elements for the JFS to perform logging. Functions relating to transaction allocation, begin, lock, and commit are managed as representations of various transactions.

RAID Data Protection

Implementing Redundant Array of Independent Disks (RAID) on the NAS appliance protects critical data and can further improve data throughput. Typical RAID levels appropriate for a NAS device include:

- *RAID Level 0*, or data striping across multiple disks, without redundancy, for optimum performance.
- *RAID Level 1*, or disk mirroring, for ultimate data protection through full redundancy.
- *RAID Level 5*, data striping with parity, for a good combination of data protection and performance.

Intel® Communications Appliance Reference Designs

Intel has developed two reference designs, both available free of charge, that provide a range of performance and fast time-to-market solutions for developers of network attached storage appliances. Intel's reference designs readily support RAID, JFS, and other software-based, value-added functionality in a compact, low-cost form factor for today's high-availability networking environments.

The Entry-Level Communications Appliance Reference Design features the Intel® Celeron™ processor, 440BX AGPset, two 82559ER 10/100 Ethernet Controllers, two IEEE 802.3 Fast Ethernet Ports, StrataFlash™ Memory/Advanced+ Boot Block flash, and support for Microsoft* Windows* 95/98/2000/NT, Linux*, Unix*, and Solaris* operating systems. The Intel Celeron processor is available at speeds of 300, 366, 433, and 566 MHz with a 66-MHz processor side bus. It includes an integrated 128-Kbyte Level 2 cache with a separate 16-Kbyte instruction and 16-Kbyte data Level 1 caches. The Level 2 cache is capable of caching 4 Gbytes of system memory address space. These components provide the processing and system-level performance headroom to meet the demands of present and future software applications.

The Value Communications Appliance Reference Design is a scalable solution that enables developers to design a single board that can be populated with an Intel Pentium® III or Celeron processor. The design features the Intel® 810 chipset, two 82559ER 10/100 Ethernet Controllers, two IDE channels (ATA/66), StrataFlash Memory/Advanced+ Boot Block flash, and support for Microsoft Windows 95/98/2000/NT, Linux, Unix, and Solaris operating systems. The Pentium III processor features an integrated, on-die, 256-Kbyte, 8-way set associative Level 2 cache and is available at 600, 700, and 850 MHz with a 100-MHz processor side bus. The processor includes a 16-Kbyte Level 1 instruction cache and 16-Kbyte Level 1 data cache. These cache arrays run at the full speed of the processor core.

Both reference designs provide numerous advantages for developers by providing a comprehensive platform solution that can dramatically accelerate time-to-market with key features important in a NAS design. With their small (5.12" x 8.5") form factor, the reference designs are ideal for compact, transportable NAS appliances. In addition, both designs support ATA/66, enabling high throughput of data. The entry-level and value reference designs provide support for various popular NAS software and RAID Levels 0, 1, and 5.

Summary

Network data storage requirements continue to grow, with no end in sight. Affordable, easy to use, compact Network Attached Storage appliances enable network users to scale storage capacity while maintaining high data availability.

The Intel® Entry-Level and Value Communications Appliance Reference Designs provide fast time-to-market development solutions for NAS appliances.

More Information

Visit Intel's Developer Web site for more information on:

- Entry-Level Communications Appliance Reference Design
- Value Communications Appliance Reference Design
- Intel Pentium III Processor
- Intel Celeron Processor
- Intel 810 chipset
- Intel® 440BX AGPset
- Scalable Performance Board Design program

Author Bio

David Hillyard directs communications platform strategy activities for Intel's Embedded Intel Architecture Division. He joined the company in 1989, and for the last several years has been involved with emerging communications services and products, including hardware, software, and system-level programs and initiatives. Before joining Intel, David served in various system-level engineering capacities at Motorola, Inc., and Digital Equipment Corp.

Software

Security and Confidentiality for Intel® Early Access Services

Umesh Shah
Technical Marketing Manager
Developer Services and Support
Intel Corporation

Overview

Intel offers software developers free access to new technology through the Intel® Early Access Services Web site. By visiting the site, developers have access to the latest, maintenance-free Intel® Itanium™-based platforms running a variety of operating systems. (See the article Access to Pre-Release IA-64 Systems in the October edition of *Intel Developer Update*.) In the past, inadequate levels of security and confidentiality have posed obstacles to the effective use of remote platforms for software development. This problem has now been resolved.

Recent major advances in security technologies have allowed Intel to resolve the issues of unauthorized access and activity. Because it is now possible to offer a high degree of security, Intel and other companies can offer shared access to remote systems, in addition to exclusive access with sys-admin privileges, for developers creating high-end software applications.

Shared and Exclusive Access

The Intel Early Access Services Web site provides developers with two levels of access:

- *Shared access* – allows several developers to share an Itanium-based system on an isolated network segment.
- *Exclusive access* – assigns an entire Itanium-based system to a single developer or to a company.

By definition, shared access over the Internet is open to multiple users. Protection from unauthorized access and actions is provided by the user's operating system (OS), the Itanium/OS platform, the virtual local-area network (VLAN), and other network security provided by the Intel data center. Shared access grants user privileges, but it does not provide developers with sys-admin privileges. Because of the moderate level of security provided by the shared access model, it is most suitable for tasks where security is not critical. These tasks can include running sample code, performing quick tests, running tutorials, transferring and executing binaries, and so on.

Exclusive access is designed to meet the requirements of developers who need robust logic resources (memory and CPU power) or who need sys-admin privileges to complete their work. For example, exclusive access allows developers to work on system software, kernel drivers, file-system drivers, and client-server development. Exclusive access grants sys-admin privileges to developers in a VLAN-isolated segment, to ensure that no abuse of those privileges occurs.

In this model, the developer (or group of developers from the same company) has full control over all system features and can use any and all of the system's available 64-bit computing capabilities. The developer does not share the system with other users, and system resources including hard disk and functionality are not restricted.

Exclusive access offers two obvious and important security benefits:

- Intel can extend better protection to the Internet connection through encryption technology.
- Full sys-admin privileges can be granted.

In addition, full control over the system allows for system-level programming, client-server development, and other activities, making an Itanium-based system virtually available on the user's private TCP/IP network.

Exclusive Access Options

Because exclusive access provides greater security, it also offers developers more options for structuring their work.

Users of exclusive access have, by default, two systems in their network segment, their local IA-32 system and the remote Itanium-based system. Developers can keep the source code on their local system, use their tools to create binaries, and then push the binaries to the remote Itanium-based system for testing. In this case, the developer's own security measures protect the source code. The disadvantage of this model is that it involves the transfer of large binary objects across the Internet connection, and this process can be very slow.

The alternative is for developers to move the source code to the IA-32 system on their assigned network segment, run the cross-platform tools there, then transfer the binaries to the Itanium-based system over the same network segment. In this case, the data center's well-designed security measures provide security and confidentiality for the source code. This technique dramatically increases the efficiency of tasks such as debugging.

When they are working with code that is highly confidential, developers at first may prefer to leave source code at their own site and push only the binaries to the remote system. As they gain greater confidence in the available level of security and confidentiality, developers can begin transferring source code to the remote system. Running and testing code on the remote 64-bit system dramatically increases productivity.

To help ensure security, exclusive-access developers can wipe all code at the end of the work period, either manually or by script, and then restore that code at the beginning of the next work period.

Security Audits

Once a security system has been designed, there are two main approaches to maintaining that security: static and live auditing. Intel uses both approaches in the early-access data center.

Static auditing involves an analysis of the overall structure of the system to identify and correct any inherent weaknesses before they can be exploited. When Intel finished designing the Itanium-based remote data center, a comprehensive static audit was performed by Internet Security Systems, a leading third-party security specialist. This organization has certified that the Intel® remote-access system appears secure from a structural perspective, including Intel's use of virtual private network (VPN), firewall, and VLAN features.

When developers use the Itanium-based systems through the network infrastructure, additional vulnerabilities may become evident. These vulnerabilities can be detected by live security auditing procedures, including dynamic and continuous monitoring of the system while it is in use. Since users access the remote Itanium-based systems via the Internet, there is a device that terminates the customer's incoming Internet connection.

At the data center, that connection is logically extended to the remote Itanium-based system through a variety of network devices, including firewalls and the Ethernet switch that manages the VLAN. Live audits of these wired connections and devices are designed to prevent unauthorized access or unwanted activities. Live auditing covers physical wire taps, in addition to detecting unauthorized and unwanted monitoring and more destructive abuse.

Intel handles physical security in a variety of ways. The data center containing the remote-access Itanium-based systems is located on Intel-owned property. The data center facility is physically separate from other work areas. It is secured behind locked doors and is subject to continuous video surveillance. Access to the data center is limited to authorized personnel only, including Intel Early-Access Services team members and system maintenance personnel.

Isolated Network Segments

To prevent users from monitoring each other's work, Intel uses a VLAN for exclusive-access users. The VLAN is implemented by a low-level Ethernet switch that creates a virtual isolated network segment for each customer, each segment with an IA-32 and Itanium-based system. The customer has exclusive use of the assigned network segment, including full sys-admin privileges for both the IA-32 and Itanium-based systems.

An isolated network segment is also used for shared-access customers. In this segment, the data center retains sys-admin privileges and assigns only user privileges to customers.

Using VLANs, two customers on separate network segments are functionally isolated from one another, with no access to other segments or systems. This means that developers cannot use their assigned IA-32 or Itanium-based systems to monitor other network segments. The VLAN also prevents a remote-access user from hacking into systems used by others.

Authentication and Legal Agreements

Intel is working to ensure that any legitimate software developer who needs access to an Itanium-based platform can obtain that access. Intel employs authentication schemes, legal agreements, and encryption to prevent unauthorized access to the system. Only customers who are approved through the application process are assigned a user name and password.

As part of the approval process, users must accept a legal agreement that defines the responsibilities and liabilities which will be accepted by the user, and by Intel, in order to maintain security and confidentiality. The user agreement specifies the level of confidentiality that will be provided. This document is available from the application area of the Intel Early Access Web site. The legal agreement will also be made available as a separate document for potential users to examine before applying for access.

Confidentiality

Confidentiality is critical for most industrial software development. Intel assures confidentiality in a variety of ways:

- Physical access to the Itanium-based data center is restricted.
- VLAN technology isolates users from each other.
- Back-ups are eliminated.
- Access is limited to users who have signed the legal agreement.

The data center is isolated from the rest of the Intel network. No Intel employee can gain access to the remote-access data center through the Intel intranet. Like any other customer, an Intel employee must go through the application process and be approved for access to the Itanium-based data center.

Just as with other customers, the work of an Intel employee is isolated in a VLAN segment. As described earlier in this article, VLAN technology effectively isolates users on shared systems. This prevents users from being able to tap into and monitor each other's work. The only people who can potentially look at customer activities are authorized data-center team members, and the personnel authorized by Intel to perform system maintenance.

Third, no back-ups are made on early-access systems either for exclusive-access or shared-access users. This restriction helps eliminate security and confidentiality concerns by avoiding duplication of data. Of course, this also means that customers are responsible for doing their own back-ups, since Intel will not have that data available for restoration purposes.

Summary

Advances in security technologies enable Intel to provide software developers with free access to secure, confidential, maintenance-free Itanium-based platforms. Developers can now use their own Internet connections to remotely access new technologies, both before and after their public release. With Intel’s extensive security measures, developers can have confidence in the use of remote platforms for important development work.

Developers can sign up now for the Intel Early Access Services, which are a feature of the Intel® Developer Services program. Self-registration can be done online at the Intel Developer’s Update site, and users can then apply to use Intel Early Access Services.

More Info

Developers can go online to learn more about the security features of Intel Early Access Services.

Intel’s legal security and confidentiality agreement can be read by following the application process. This agreement will also be made available soon through the main Intel Early Access Services Web page.

The early-access program is a feature of Intel Developer Services.

Author Bio

Umesh Shah has been with Intel for nine years. He joined Intel as a VLSI chip designer for the i360SL processor and was a member of a design team for i486SL processor. He then worked in software design for the Mobile Companion (PDA). For seven years Umesh has focused on Pentium® processor projects, including real-time address tracing, the Instruction Tracer for the Pentium and P6 microarchitecture, the event-based profiler, and enabling technologies for Itanium-based server software. Currently, Umesh is a technical marketing manager with IAG/SEG/Developer Services and project manager for Intel Early Access Services. He received his B.S. in electrical engineering from BITS, Pilani, India and his M.S. in computer engineering from Arizona State University.

—End of Intel Developer Update Magazine Issue 14—